

# EXHIBIT G



**CISA**  
CYBER+INFRASTRUCTURE

TLP:WHITE

# Alert (TA13-088A)

[More Alerts](#)

## DNS Amplification Attacks

Original release date: March 29, 2013 | Last revised: June 04, 2019

### Systems Affected

- Domain Name System (DNS) servers

### Overview

A Domain Name Server (DNS) amplification attack is a popular form of distributed denial of service (DDoS) that relies on the use of publically accessible open DNS servers to overwhelm a victim system with DNS response traffic.

### Description

A Domain Name Server (DNS) Amplification attack is a popular form of Distributed Denial of Service (DDoS), in which attackers use publically accessible open DNS servers to flood a target system with DNS response traffic. The primary technique consists of an attacker sending a DNS name lookup request to an open DNS server with the source address spoofed to be the target's address. When the DNS server sends the DNS record response, it is sent instead to the target. Attackers will typically submit a request for as much zone information as possible to maximize the amplification effect. In most attacks of this type observed by US-CERT, the spoofed queries sent by the attacker are of the type, "ANY," which returns all known information about a DNS zone in a single request. Because the size of the response is considerably larger than the request, the attacker is able to increase the amount of traffic directed at the victim. By leveraging a botnet to produce a large number of spoofed DNS queries, an attacker can create an immense amount of traffic with little effort. Additionally, because the

TLP:WHITE

responses are legitimate data coming from valid servers, it is extremely difficult to prevent these types of attacks. While the attacks are difficult to stop, network operators can apply several possible mitigation strategies.

While the most common form of this attack that US-CERT has observed involves DNS servers configured to allow unrestricted recursive resolution for any client on the Internet, attacks can also involve authoritative name servers that do not provide recursive resolution. The attack method is similar to open recursive resolvers, but is more difficult to mitigate since even a server configured with best practices can still be used in an attack. In the case of authoritative servers, mitigation should focus on using Response Rate Limiting to restrict the amount of traffic.

## Impact

A misconfigured Domain Name System (DNS) server can be exploited to participate in a distributed denial of service (DDoS) attack.

## Solution

## DETECTION

While it is not easy to identify authoritative name servers used in DNS reflection attacks as vulnerability is not caused by a misconfiguration, there are several freely available options for detecting open recursive resolvers. Several organizations offer free, web-based scanning tools that will search a network for vulnerable open DNS resolvers. These tools will scan entire network ranges and list the address of any identified open resolvers.

### *Open DNS Resolver Project*

<http://openresolverproject.org>

The Open DNS Resolver Project has compiled a list of DNS servers that are known to serve as globally accessible open resolvers. The query interface allows network administrators to enter IP ranges in CIDR format [1].

### *The Measurement Factory*

<http://dns.measurement-factory.com>

Like the Open DNS Resolver Project, the Measurement Factory maintains a list of Internet accessible DNS servers and allows administrators to search for open recursive resolvers [2]. In addition, the Measurement Factory offers a free tool to test a single DNS resolver to determine if it allows open recursion. This will allow an administrator to determine if configuration

changes are required and verify that configuration changes have been successful [3]. Finally, the site offers statistics showing the number of public resolvers detected on the different Autonomous System (AS) networks, sorted by the highest number found [4].

### *DNSInspect*

<http://www.dnsinspect.com>

Another freely available, web-based tool for testing DNS resolvers is DNSInspect. This site is similar to The Measurement Factory's ability to assess an individual resolver for vulnerability, but offers the ability to test an entire DNS Zone for several other possible configuration and security issues [5].

### Indicators

In a typical recursive DNS query, a client sends a query request to a local DNS server requesting the resolution of a name or the reverse resolution of an IP address. The DNS server performs the necessary queries on behalf of the client and returns a response packet with the requested information or an error [6, page 21]. The specification does not allow for unsolicited responses. In a DNS amplification attack, the main indicator is a query response without a matching request.

## MITIGATION

Unfortunately, due to the massive traffic volume that can be produced by one of these attacks, there is often little that the victim can do to counter a large-scale DNS amplification-based distributed denial-of-service attack. However, it is possible to reduce the number of servers that can be used by attackers to generate the traffic volumes.

While the only effective means of eliminating the use of recursive resolvers in this type of attack is to eliminate unsecured recursive resolvers, this requires an extensive effort by various parties. According to the Open DNS Resolver Project, of the 27 million known DNS resolvers on the Internet, approximately "25 million pose a significant threat" of being used in an attack [1]. However, several possible techniques are available to reduce the overall effectiveness of such attacks to the Internet community as a whole. Where possible, configuration links have been provided to assist administrators with making the recommended changes. The configuration information has been

limited to BIND9 and Microsoft's DNS Server, which are two widely deployed DNS servers on federal networks. If you are running a different DNS server, please consult your vendor's documentation for configuration details.

TLP:WHITE

## Source IP Verification

Because the DNS queries being sent by the attacker-controlled clients must have a source address spoofed to appear as the victim's system, the first step to reducing the effectiveness of DNS amplification is for Internet Service Providers to reject any DNS traffic with spoofed addresses. The Network Working Group of the Internet Engineering Task Force released Best Current Practice 38 document in May 2000 and Best Current Practice 84 in March 2004 that describes how an Internet Service Provider can filter network traffic on their network to reject packets with source addresses not reachable via the actual packet's path [7]. The changes recommended in this document would cause a routing device to evaluate whether it is possible to reach the source address of the packet via the interface that transmitted the packet. If it is not possible, then the packet obviously has a spoofed source address. This configuration change would substantially reduce the potential for most popular types of DDoS attacks. As such, we highly recommend to all network operators to perform network ingress filtering if possible.

## Disabling Recursion on Authoritative Name Servers

Many of the DNS servers currently deployed on the Internet are exclusively intended to provide name resolution for a single domain. In these systems, DNS resolution for private client systems may be provided by a separate server and the authoritative server acts only as a DNS source of zone information to external clients. These systems do not need to support recursive resolution of other domains on behalf of a client, and should be configured with recursion disabled.

### Bind9

Add the following to the global options [8]:

```
options {  
    allow-query-cache { none; };  
    recursion no;  
};
```

### Microsoft DNS Server

TLP:WHITE

In the Microsoft DNS console tool [9]:

TLP:WHITE

1. Right-click the DNS server and click Properties.
2. Click the Advanced tab.
3. In Server options, select the "Disable recursion" check box, and then click OK.

## Limiting Recursion to Authorized Clients

For DNS servers that are deployed within an organization or Internet Service Provider, the resolver should be configured to perform recursive queries on behalf of authorized clients only. These requests typically should only come from clients within the organization's network address range. We highly recommend that all server administrators restrict recursion to only clients on the organization's network.

### BIND9

In the global options, include the following [10]:

```
acl corpnets { 192.168.1.0/24; 192.168.2.0/24; };
options {
    allow-query { any; };
    allow-recursion { corpnets; };
};
```

### Microsoft DNS Server

It is not currently possible to restrict recursive DNS requests to a particular client address range in Microsoft DNS Server. To approximate the functionality of the BIND access control lists in Microsoft's DNS Server, a different caching-only name server should be set up internally to provide recursive resolution. A firewall rule should be created to block incoming access to the caching-only server from outside the organization's network. The authoritative name server functionality would then need to be hosted on a separate server, but configured to disable recursion as previously described.

## Response Rate Limiting (RRL)

There is currently an experimental feature available as a set of patches for BIND9 that allows an administrator to limit the maximum number of responses per second being sent to one client from the name server [11]. This functionality is intended to be used on authoritative domain name servers only as it will affect performance on recursive resolvers. To provide

TLP:WHITE

the most effective protection, we recommend that authoritative and recursive name servers run on different systems, with RRL implemented on the authoritative server and access control lists implemented on the recursive server. This will reduce the effectiveness of DNS amplification attacks by reducing the amount of traffic coming from any single authoritative server while not affecting the performance of the internal recursive resolvers.

TLP:WHITE

## BIND9

There are currently patches available for 9.8.latest and 9.9.latest to support RRL on UNIX systems. Red Hat has made updated packages available for Red Hat Enterprise Linux 6 to provide the necessary changes in advisory RHSA-2013:0550-1. On BIND9 implementation running the RRL patches, include the following lines to the options block of the authoritative views [12]:

```
rate-limit {  
    responses-per-second 5;  
    window 5;  
};
```

## Microsoft DNS Server

In Windows Server 2016, the `Set-DnsServerResponseRateLimiting` cmdlet enables RRL with default settings[15][16]. See more settings at `Set-DnsServerResponseRateLimiting`.

**Disclaimer:** Rate limiting DNS responses may prevent legitimate hosts from receiving answers. Such hosts may be at increased risk for successful DNS cache poisoning attacks.

RRL of DNS responses may prevent legitimate hosts from receiving answers. Such hosts may be at increased risk for successful DNS cache poisoning attacks.

## References

- [1] Open DNS Resolver Project
- [2] The Measurement Factory, "List Open Resolvers on Your Network"
- [3] The Measurement Factory, "Open Resolver Test"
- [4] The Measurement Factory, "Open Resolvers for Each Autonomous System"
- [5] "DNSInspect," DNSInspect.com
- [6] RFC 1034: DOMAIN NAMES - CONCEPTS AND FACILITIES
- [7] BCP 38: Network Ingress Filtering: Defeating Denial of Service Attacks whic...
- [8] Chapter 3. Name Server Configuration
- [9] Disable recursion on the DNS server
- [10] Chapter 7. BIND 9 Security Considerations
- [11] DNS Response Rate Limiting (DNS RRL)
- [12] Response Rate Limiting in the Domain Name System (DNS RRL)

TLP:WHITE

- [13] The Measurement Factory, "Open Resolvers for Each Autonomous System"
- [14] Configure a DNS Server to Use Forwarders
- [15] Response Rate Limiting in Windows DNS Server
- [16] What's New in DNS Server in Windows Server

TLP:WHITE

## Revisions

March 29, 2013: Initial release

April 18, 2013: Minor updates to Description and Solution sections(Source IP Verification and BIND9)

July 5, 2013: Added disclaimer for DNS request rate limiting

July 8, 2013: Updates to Description, Detection, and Mitigation sections

July 22, 2013: Minor updates to recursion and RRL advice

June 4, 2019: Added Windows DNS Server support for RRL

**This product is provided subject to this Notification and this Privacy & Use policy.**

TLP:WHITE





**CISA**  
CYBER+INFRASTRUCTURE

TLP:WHITE

# Alert (TA14-013A)

[More Alerts](#)

## NTP Amplification Attacks Using CVE-2013-5211

Original release date: January 13, 2014 | Last revised: October 06, 2016

### Systems Affected

NTP servers

### Overview

A Network Time Protocol (NTP) Amplification attack is an emerging form of Distributed Denial of Service (DDoS) that relies on the use of publically accessible NTP servers to overwhelm a victim system with UDP traffic.

### Description

The NTP service supports a monitoring service that allows administrators to query the server for traffic counts of connected clients. This information is provided via the "monlist" command. The basic attack technique consists of an attacker sending a "get monlist" request to a vulnerable NTP server, with the source address spoofed to be the victim's address.

### Impact

The attack relies on the exploitation of the 'monlist' feature of NTP, as described in CVE-2013-5211, which is enabled by default on older NTP-capable devices. This command causes a list of the last 600 IP addresses which connected to the NTP server to be sent to the victim. Due to the spoofed source address, when the NTP server sends the response it is sent instead to the victim. Because the size of the response is typically considerably larger than the request, the attacker is able to amplify the volume of traffic directed at the victim. Additionally, because the responses are legitimate data coming from valid servers, it is especially difficult to block

TLP:WHITE

these types of attacks. The solution is to disable “monlist” within the NTP server or to upgrade to the latest version of NTP (4.2.7) which disables the “monlist” functionality.

## Solution

### Detection

On a UNIX-platform, the command “ntpd” will query existing NTP servers for monitoring data. If the system is vulnerable to exploitation, it will respond to the “monlist” command in interactive mode. By default, most modern UNIX and Linux distributions allow this command to be used from localhost, but not from a remote host. To test for monlist support, execute the following command at the command line:

```
/usr/sbin/ntpd <remote server>  
  
monlist
```

Additionally, the “ntp-monlist” script is available for NMap, which will automatically display the results of the monlist command. If the system does not support the monitor query, and is therefore not vulnerable to this attack type, NMap will return an error type 4 (No Data Available) or no reply at all.

### Recommended Course of Action

As all versions of ntpd prior to 4.2.7 are vulnerable by default, the simplest recommended course of action is to upgrade all versions of ntpd that are publically accessible to at least 4.2.7. However, in cases where it is not possible to upgrade the version of the service, it is possible to disable the monitor functionality in earlier versions of the software.

To disable “monlist” functionality on a public-facing NTP server that cannot be updated to 4.2.7, add the “noquery” directive to the “restrict default” line in the system’s ntp.conf, as shown below:

```
restrict default kod nomodify notrap nopeer  
noquery  
  
restrict -6 default kod nomodify notrap nopeer  
noquery
```

## References

## Revisions

January 13, 2014 - Initial Release

**This product is provided subject to this Notification and this Privacy & Use policy.**



# Alert (TA14-017A)

[More Alerts](#)

## UDP-Based Amplification Attacks

Original release date: January 17, 2014 | Last revised: March 02, 2018

### Systems Affected

Certain application-layer protocols that rely on the User Datagram Protocol (UDP) have been identified as potential attack vectors. These include

- Domain Name System (DNS),
- Network Time Protocol (NTP),
- Connection-less Lightweight Directory Access Protocol (CLDAP),
- Character Generator Protocol (CharGEN),
- Simple Service Discovery Protocol (SSDP),
- BitTorrent,
- Simple Network Management Protocol version 2 (SNMPv2),
- Kad,
- Portmap/Remote Procedure Call (RPC),
- Quote of the Day (QOTD),
- Multicast Domain Name System (mDNS),
- Network Basic Input/Output System (NetBIOS),
- Quake Network Protocol,
- Steam Protocol,
- Routing Information Protocol version 1 (RIPv1),
- Lightweight Directory Access Protocol (LDAP),
- Trivial File Transfer Protocol (TFTP), and
- Memcached.

### Overview

A distributed reflective denial-of-service (DRDoS) is a form of distributed denial-of-service (DDoS) attack that relies on publicly accessible UDP servers and bandwidth amplification factors (BAFs) to overwhelm a victim's system with UDP traffic.

## Description

TLP:WHITE

By design, UDP is a connection-less protocol that does not validate source Internet Protocol (IP) addresses. Unless the application-layer protocol uses countermeasures such as session initiation in Voice over Internet Protocol, an attacker can easily forge the IP packet datagram (a basic transfer unit associated with a packet-switched network) to include an arbitrary source IP address. [1] When many UDP packets have their source IP address forged to the victim IP address, the destination server (or amplifier) responds to the victim (instead of the attacker), creating a reflected denial-of-service (DoS) attack.

Certain commands to UDP protocols elicit responses that are much larger than the initial request. Previously, attackers were limited by the linear number of packets directly sent to the target to conduct a DoS attack; now a single packet can generate between 10 and 100 times the original bandwidth. This is called an amplification attack, and when combined with a reflective DoS attack on a large scale, using multiple amplifiers and targeting a single victim, DDoS attacks can be conducted with relative ease.

The potential effect of an amplification attack can be measured by BAF, which can be calculated as the number of UDP payload bytes that an amplifier sends to answer a request, compared to the number of UDP payload bytes of the request. [2] [3]

The following is a list of known protocols and their associated BAFs. US-CERT offers thanks to Christian Rossow for providing this information. For more information on BAFs, please see Christian's blog and associated research paper.

Protocol	Bandwidth Amplification Factor	Vulnerable Command
DNS	28 to 54	see: TA13-088A [4]
NTP	556.9	see: TA14-013A [5]
SNMPv2	6.3	GetBulk request
NetBIOS	3.8	Name resolution
SSDP	30.8	SEARCH request
CharGEN	358.8	Character generation request
QOTD	140.3	Quote request
BitTorrent	3.8	File search
Kad	16.3	Peer list exchange
Quake Network Protocol	63.9	Server info exchange
Steam Protocol	5.5	Server info exchange
Multicast DNS (mDNS)	2 to 10	Unicast query
RIPv1	131.24	Malformed request
Portmap (RPCbind)	7 to 28	Malformed request
LDAP	46 to 55	Malformed request [6]

TLP:WHITE

CLDAP [7]	56 to 70	—	TLP:WHITE
TFTP [23]	60	—	
Memcached [25]	10,000 to 51,000	—	

In March 2015, the CERT Coordination Center of the Software Engineering Institute issued Vulnerability Note VU#550620 describing the use of mDNS in DRDoS attacks. Attackers can leverage mDNS by sending more information than can be handled by the device, thereby causing a DoS condition. [8]

In July 2015, Akamai Technologies' Prolexic Security Engineering and Research Team (PLXsert) issued a threat advisory describing a surge in DRDoS attacks using RIPv1. Malicious actors are leveraging the behavior of RIPv1 for DDoS reflection through specially crafted request queries. [9]

In August 2015, Level 3 Threat Research Labs reported a new form of DRDoS attack that uses portmap. Attackers are leveraging the behavior of the portmap service through spoofed requests to flood a victim's network with UDP traffic. [10]

In October 2016, Corero Network Security reported a new DDoS amplification attack exploiting LDAP directory services servers against its customers. [11]

In November 2017, Netlab 360 reported that CLDAP is now the third most common DRDoS attack, behind DNS and NTP attacks. [12]

In February 2018, SENKI reported an increase in Memcached-based reflection DDoS attacks (via UDP/TCP port 11211) with an unprecedented amplification factor. [24]

## Impact

Attackers can use the bandwidth and relative trust of large servers that provide the UDP protocols provided in this alert to flood victims with unwanted traffic and create a DDoS attack.

## Solution

### Detection

Detection of DRDoS attacks is not easy because of their use of large, trusted servers that provide UDP services. Network operators of these exploitable services may apply traditional DoS mitigation techniques. To detect a DRDoS attack, watch out for abnormally large responses to a particular IP address, which may indicate that an attacker is using the service.

TLP:WHITE

There are a few things victims of DRDoS attacks can do to detect such activity and respond:

1. Detect and alert large UDP packets to higher order ports.
2. Detect and alert on any non-stateful UDP packets. (A simple Snort example is below. The approach will need to be customized to each environment with a whitelist and known services.)

Simple Snort rule example for stateless UDP check

```
var HOME_NET [10.10.10.20]
preprocessor stream5_global: track_ip yes, track_tcp yes, track_udp
yes, track_icmp no, max_tcp 262144, max_udp 131072
preprocessor stream5_ip: timeout 180
preprocessor stream5_tcp: policy first, use_static_footprint_sizes
preprocessor stream5_udp: timeout 180, ignore_any_rules
alert udp HOME_NET 1024: -> any any (msg:"UDP Session start";
flowbits:set,logged_in; flowbits:noalert; sid: 1001;)
alert udp any any -> HOME_NET 1024: (msg:"UDP Stateless";
flowbits:isnotset,logged_in; sid: 1002)
```

3. Upstream providers should maintain updated contacts and methods with downstream customers to send alerts by network.

In general, network and server administrators for Internet service providers (ISPs) should use the following best practices to avoid becoming amplifier nodes:

1. Use network flow to detect spoofed packets. (See the Mitigation section below for information on verifying spoofed traffic before blocking that traffic.)
2. Use network flow or other summarized network data to monitor for an unusual number of requests to at-risk UDP services.
3. Use network flow to detect service anomalies (e.g., bytes-per-packet and packets-per-second anomalies).

## Mitigation

The following steps can help mitigate a DRDoS attack:

1. Use stateful UDP inspections—such as reflexive access control lists—to reduce the impact to critical services on border firewalls or border routers.
- [13]

2. Use a Border Gateway Protocol (BGP) to create a Remotely Triggered Blackhole, preferably in coordination with upstream providers or ISPs. [14]
3. Maintain a list of primary upstream provider emergency contacts to coordinate responses to attacks. Upstream providers should conduct mitigation in coordination with downstream customers.

In general, ISP network and server administrators should use the following best practices to avoid becoming amplifier nodes:

1. Regularly update software and configurations to deny or limit abuse (e.g., DNS response rate limit). [15] [16] [17]
2. Disable and remove unwanted services, or deny access to local services over the Internet.
3. Use UDP-based protocols—e.g., quality of service (QoS) on switching and routing devices—to enable network-based rate-limiting to legitimate services provided over the Internet.
4. Work with Customer Provider Edge manufacturers for secure configuration and software. [18]

As a service provider, to avoid any misuse of Internet resources:

1. Use ingress filtering to block spoofed packets (See the Spoofer Project [19], and IETF BCP 38 and BCP 84 guidelines). [20]
2. Use traffic shaping on UDP service requests to ensure repeated access to over-the-Internet resources is not abusive. [21] [22]

## References

- [1] SIP: Session Initiation Protocol
- [2] Amplification Hell: Abusing Network Protocols for DDoS
- [3] Amplification Hell: Revisiting Network Protocols for DDoS Abuse
- [4] DNS Amplification Attacks
- [5] NTP Amplification Attacks Using CVE-2013-5211
- [6] Open LDAP Scanning Project
- [7] CLDAP Reflection DDoS
- [8] VU#550620: Multicast DNS (mDNS) implementations may respond to unicast quer...
- [9] RIPv1 Reflection DDoS [Medium Risk]
- [10] A New DDoS Reflection Attack: Portmapper; An Early Warning to the Industry
- [11] Corero Warns of Powerful New DDoS Attack Vector
- [12] CLDAP is Now the No.3 Reflection Amplified DDoS Attack Vector, Surpassing ...
- [13] Configuring IP Session Filtering (Reflexive Access Lists)
- [14] Remotely-Triggered Black Hole (RTBH) Routing
- [15] A Quick Introduction to Response Rate Limiting
- [16] Network Ingress Filtering: Defeating Denial of Service Attacks Which Emplo...
- [17] Ingress Filtering for Multihomed Networks
- [18] Abuse of Customer Premise Equipment and Recommended Actions



- [19] The Spoofer Project
- [20] Abuse of Customer Premise Equipment and Recommended Actions
- [21] An Architecture for Differentiated Services
- [22] New Terminology and Clarifications for Diffserv
- [23] TFTP DDoS Amplification Attack
- [24] Memcached on Port 11211 UDP & TCP Being Exploited
- [25] Open Memcached Key-Value Store Scanning Project

TLP:WHITE

## Revisions

February 9, 2014 – Initial Release

March 7, 2014 – Updated page to include research links

July 13, 2015 – Added RIPv1 as an attack vector

August 19, 2015 – Added Multicast DNS (mDNS) and Portmap (RPCbind) as attack vectors

April 13, 2016 – Updated detection and mitigation information

November 4, 2016 – Updated for LDAP attack vector

December 4, 2017 – Added information on CLDAP as an attack vector

December 6, 2017 – Added information on TFTP as an attack vector

February 27, 2018 – Added information on Memcached as an attack vector

March 2, 2018 - Updated information on Memcached as an attack vector

**This product is provided subject to this Notification and this Privacy & Use policy.**

TLP:WHITE